Aerodynamic Shape Optimization of Wind Turbine Blades Using a Reynolds-Averaged Navier–Stokes Model and an Adjoint Method

Tristan Dhert, Turaj Ashuri, and Joaquim R. R. A. Martins

Abstract Computational fluid dynamics (CFD) is increasingly used to analyze wind turbines, and the next logical step is to develop CFD-based optimization to enable further gains in performance and reduce model uncertainties. We present an aerodynamic shape optimization framework consisting of a Reynolds-averaged Navier–Stokes solver coupled to a numerical optimization algorithm, a geometry modeler, and a mesh perturbation algorithm. To efficiently handle the large number of design variables, we use a gradient-based optimization technique together with an adjoint method for computing the gradients of the torque coefficient with respect to the design variables. To demonstrate the effectiveness of the proposed approach, we maximize the torque of the NREL VI wind turbine blade with respect to pitch, twist, and airfoil shape design variables while constraining the blade thickness. We present a series of optimization cases with increasing number of variables, both for a single wind speed and for multiple wind speeds. For the optimization at a single wind speed performed with respect to all the design variables (1 pitch, 11 twist, and 240 airfoil shape variables), the torque coefficient increased by 22.4% relative to the NREL VI design. For the multiple-speed optimization, the torque increased by an average of 22.1%. Depending on the CFD mesh size and number of design variables, the optimization time ranges from 2 to 24 h when using 256 cores, which means that wind turbine designers can use this process routinely.

Keywords: wind turbine design, aerodynamic shape optimization, computational fluid dynamics, NREL VI, numerical optimization, adjoint methods

1 Introduction

Wind energy has experienced a considerable reduction in the levelized cost of energy (LCOE) over the last few years. However, further cost reduction is needed for wind energy to become more competitive with traditional energy resources such as coal and natural gas. Besides decreasing the turbine capital cost (TCC) and balance of station, the annual energy production (AEP) has to increase to further reduce costs. This reduction can be achieved by upscaling wind turbines, developing new optimization methods to further optimize wind turbine designs, and introducing new wind turbine concepts and components.

The complexity of this challenge increases with the size of wind turbines [1], and placing wind turbines offshore further exacerbates this problem [2]. Given the multidisciplinary nature of wind turbines, their design hinges on optimizing the trade-off between aerodynamic performance, structural efficiency, and operational and manufacturing cost [3]. Therefore, multidisciplinary design optimization (MDO) [4] could become an important tool in the design of the next generation of

1

high-performance wind turbines by maximizing the LCOE while accounting for all relevant disciplines and interactions [5].

Previous studies have tackled several areas in wind turbine design, ranging from the optimization of rotor design [6, 7, 8, 9, 10] to support-structure design [11, 12, 13]. These studies used various levels of fidelity in the models, ranging from low-fidelity (blade element momentum-based fidelity) [14, 15, 16, 17, 18, 19, 20] to higher-fidelity (free and prescribed vortex-based fidelity) [21, 22, 23, 24, 25, 26, 27].

Currently, most wind turbine blade designs are based on low-fidelity models because of their ease of implementation, lower computational cost, and fast convergence to feasible solutions. However, these models do not accurately represent compressibility, viscosity, and the three-dimensional (3D) effects present in wind turbines, which results in suboptimal designs. These important effects can only be captured accurately by using 3D computational fluid dynamics (CFD) that includes viscous terms.

One approach adopted to incorporate high-fidelity simulations into wind turbine design is to use two-dimensional Reynolds-averaged Navier–Stokes (RANS) CFD solvers to optimize wind turbine airfoils, neglecting 3D effects [28, 29]. Another effort considered 3D CFD to optimize turbine blade winglets with respect to two design variables [30]. However, CFD-based aerodynamic shape optimization for wind turbine blades has not yet been done because of the compounding challenges of high-cost simulations and large numbers of shape design variables. In recent work, Economon et al. [31] used a RANS solver and a continuous adjoint approach to optimize the NREL VI wind turbine blade with respect to 50 airfoils shape variables. Their optimization resulted in a 4% increase in torque, but only three design optimization iterations were completed. Large-eddy simulation (LES) provide even higher fidelity than RANS for massively separated flow conditions, but the computational cost is prohibitive for inclusion in an optimization cycle.

These challenges have been addressed in the aerospace engineering community, where CFD-based design optimization has been used for decades, with applications to wing design [32, 33, 34, 35] and helicopter blade design [36, 37, 38, 36, 39, 40]. Given that wind turbine blades are in essence rotating wings, and share some of the physics in helicopter blades, much of the design optimization techniques developed by the aerospace community can be adapted to the design optimization of wind turbine blades.

In the present work, we address the challenges of CFD-based wind turbine optimization by developing a robust and efficient aerodynamic shape optimization framework. We use an arbitrary Lagrangian–Eulerian RANS model for the CFD analysis and a discrete adjoint method that efficiently computes the gradients of the performance metrics with respect to the design variables. In contrast to lower-fidelity models, the RANS CFD used herein captures the complex flow phenomena present in wind turbines, resulting in a more reliable performance analysis and ultimately a more refined design.

Unsteady effects such as tower shadow and dynamic stall are neglected in this work, since performing design optimization based on unsteady CFD is currently too costly. While we perform some verification and validation of our CFD model against the NREL VI experimental results, the focus of this work is on demonstrating the first CFD-based wind turbine optimization that considers a comprehensive set of airfoil shape design variables. The validation of CFD models, including the validity of turbulence models, is still the subject of ongoing research in the wind energy community.

To demonstrate the proposed approach, we present a series of torque coefficient maximization results starting with the NREL VI as the baseline design. These optimizations range from a single wind speed case with pitch as a design variable to a case that optimizes 251 design variables for multiple wind speeds within the rated region.

The remainder of this paper is organized as follows: First, we present the aerodynamic shape

optimization methodology for rotating flow problems. Next, we apply the methodology to the NREL VI wind turbine blade as a benchmark case [41]. This requires a problem statement fitted for this particular case. Finally, we present the results of the NREL VI wind turbine blade optimization, followed by the conclusions.

2 Methodology for aerodynamic shape optimization

We now describe the various aerodynamic shape optimization framework components used in this work. In addition to the CFD solver, the following components are required: geometry parametrization, mesh deformation, gradient computation, and numerical optimization algorithm.

2.1 Geometric parametrization

To optimize a given shape, we need to parametrize it by using a set of shape parameters that the optimizer ultimately controls. This parametrization must result in smooth shape changes to ensure that our gradient-based optimizer is effective. In addition, the geometry parametrization should be differentiable, and the gradients of the surface coordinates with respect to the shape parameters should be computed accurately and efficiently.

In this work, we use the free-form deformation (FFD) volume approach to parametrize the blade geometry by embedding the design inside a hexahedron [42]. This approach is known to have an efficient and compact set of geometry design variables, which facilitates the manipulation of complex geometries. Once the local coordinates of the vertices of the geometry are expressed in the FFD volume by doing a Newton search, FFD control points on the surface of the volume are used to deform the baseline design. When the lattice structure within the volume deforms by moving the surface control points, the FFD method deforms the embedded object as well. We use an FFD volume composed of B-spline tensor patches in three dimensions; this is called a trivariate B-spline volume. B-spline volumes are chosen because derivatives of any point inside the volume are easily computed [43].

Two FFD volumes are constructed for the two-bladed NREL VI rotor: one for each blade. The control-point parameters of the two FFD volumes are linked to ensure that the two blades have exactly the same shape. Figure 1 shows the FFD volumes for the original NREL VI rotor. The control points are fixed in the cylindrical piece near the root because this section is not shaped for aerodynamic performance.

Although the basic shape variables correspond to the movement of the control points at the surface of the FFD volume, we also implement composite shape variables corresponding global shape changes, such as blade pitch and the spanwise variation of twist, which rotate the blade sections about its 25% chord line. The twist can change in the 11 outboard sections (the root section is set to match the blade pitch). The spanwise interpolation of the twist distribution is done by using a cubic spline within the FFD. The airfoil is shaped at the 12 outboard sections (i.e., the whole blade excluding the cylindrical piece next to the hub). Each airfoil section is shaped by moving 20 FFD control points (10 on the lower surface and 10 on the upper surface), resulting in a total of 240 airfoil shape variables. Our previous work on wing design shows that this number of shape-design variables provides a good trade-off between computational effort and quality of the optimum [35]. The geometry parametrization can also handle chord and span variation, but we do not exploit them in this work because this would require introducing other disciplines, such as structures, to make the appropriate optimization trade-offs [44].
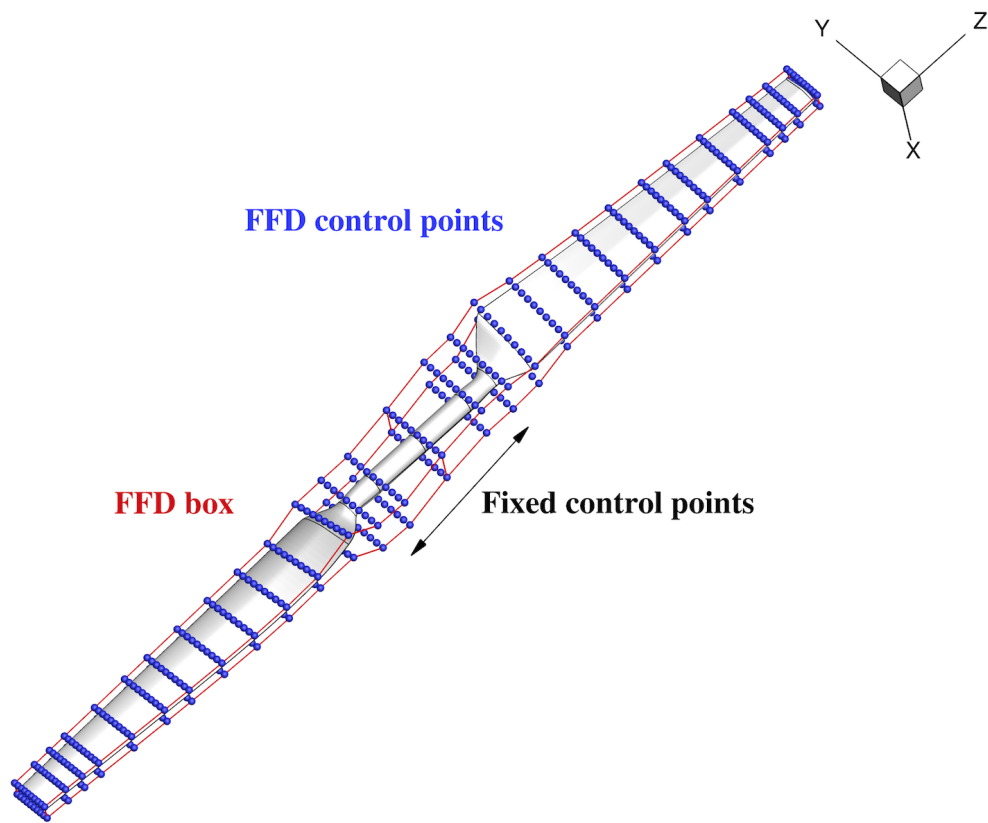
Figure 1: Parametrization of FFD volume shape for the NREL VI rotor.

2.2 Mesh deformation

Once a new blade surface shape is determined from the set of shape variables obtained by using FFD, surface deformations are applied to the CFD surface mesh, and then propagated to the entire CFD volume mesh.

   To propagate the surface-shape deformations into the volume mesh, we use the hybrid approach developed by Kenway et al. [43], which combines algebraic and linear elasticity methods. The main idea of the hybrid approach is to apply a linear elasticity method on a coarser approximation of the mesh for large, low-frequency perturbations, while using an algebraic approach to attenuate small, high-frequency deformations. The hybrid approach consists of the following steps:

1. Select a subset of nodes for each mesh edge. These nodes form a coarser approximation of the complete mesh.

2. Deform this coarser approximation of the initial mesh by using the linear elasticity method.

3. Regenerate each block with linear or cubic-Hermite spine interpolation to obtain the full mesh approximation.

4. Find the remainder of the deformed mesh by using the algebraic technique.

As an example, we compare the original and deformed mesh of the NREL VI S809 airfoil in Figure 2.

(a) Original S809 airfoil.　　　　　　　　　　　(b) Deformed S809 airfoil.
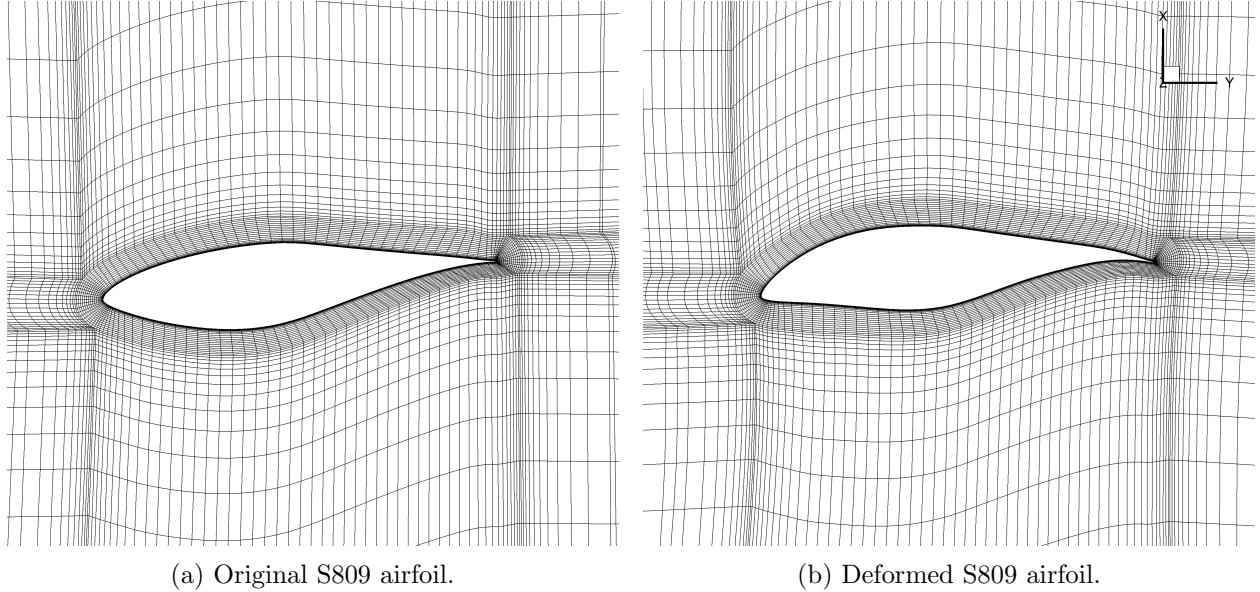
Figure 2: Mesh deformation of the NREL VI blade at $z = 3.3$ m.

2.3 Computational fluid dynamics solver

In this research, we use the RANS-based CFD solver ADflow, which is an extension of the SUmb solver [45] that adds adjoint and optimization capability [46]. ADflow uses parallel computing to solve the Euler, laminar Navier–Stokes, and RANS equations in either steady, unsteady, or time-spectral modes [47]. Depending on the size of the computational domain, a cluster of 16 to 256 processors is used for the cases presented in this paper. ADflow is a finite-volume, cell-centered multi-block CFD solver that provides options for a variety of turbulence models with one, two, and four equations, and a variety of adaptive wall functions. The Roe scheme [48] with no limiter and

a four stage Runge–Kutta scheme [49] with a segregated SA turbulence model [50] are used in the design optimizations presented herein. It is a common practice to use the k-ϵ turbulence model for wind-energy applications; however, because we are simulating the rotor in the attached flow conditions, the SA and k-ϵ turbulence models yield similar results.

When considering an isolated rotor rotating at a constant angular velocity $\boldsymbol{\omega} = [\omega_x, \omega_y, \omega_z]^T$, where the rotor plane is perpendicular to a uniform steady wind in a fixed reference frame, the flow field is intrinsically unsteady. However, to simplify the CFD analysis and make design optimization tractable, we approximate the flow as a steady-state solution by solving the problem in a reference frame co-rotating with the rotor. Unsteady flow can still exist in the presence of unsteady phenomena, such as flow separation. By using this approach, the computational domain does not rotate, but the flow field incorporates a rotational term. Assuming that both reference frames have the same origin, the velocity $\boldsymbol{u}_\omega$ of the co-rotating reference frame can be related to the velocity $\boldsymbol{u}$ in a fixed reference frame as follows:

$$\boldsymbol{u} = \boldsymbol{u}_\omega + \boldsymbol{\omega} \times \boldsymbol{x}, \tag{1}$$

where $\boldsymbol{x} = [x - x_0, y - y_0, z - z_0]^T$ is the position vector pointing from the origin of the reference frame (x_0, y_0, z_0) to a point (x, y, z) in the flow domain. The time derivative of a time-dependent vector function $\boldsymbol{f}(t)$ in the co-rotating reference frame is

$$\frac{d\boldsymbol{f}(t)}{dt} = \left(\frac{d\boldsymbol{f}(t)}{dt}\right)_\omega + \boldsymbol{\omega} \times \boldsymbol{f}(t). \tag{2}$$

The conservation of momentum of an arbitrary time-independent control volume Ω with permeable boundary $\partial\Omega$ is governed by the Navier–Stokes equations. By using equations (1) and (2) in the momentum equations for a fixed reference frame, conservation of momentum can be expressed in the integral conservation form with co-rotating motion (assuming no external forces):

$$\frac{\partial}{\partial t} \int_\Omega \rho \boldsymbol{u} d\Omega + \int_{\partial\Omega} \rho \boldsymbol{u} \left([\boldsymbol{u} - (\boldsymbol{\omega} \times \boldsymbol{x})]^T \cdot \boldsymbol{n}\right) dS = -\int_{\partial\Omega} (p\boldsymbol{n} - \underline{\tau} \cdot \boldsymbol{n}) dS - \int_\Omega \rho (\boldsymbol{\omega} \times \boldsymbol{u}) d\Omega, \tag{3}$$

where ρ is the local density, p is the static pressure, $\underline{\tau}$ is the viscous stress tensor (assuming a Newtonian fluid), $\boldsymbol{u} = [u, v, w]^T$ is the local velocity of the fluid, and $\boldsymbol{n} = [n_x, n_y, n_z]^T$ is the outward-pointing unit normal vector. When using the co-rotating reference frame, the momentum equations gain an extra body force that represents the combined effect of the centrifugal and Coriolis forces. Since the forces are perpendicular to the direction of motion, no work is produced and thus no additional term appears in the energy equation.

2.4   Adjoint formulation

As previously mentioned, RANS-based aerodynamic shape optimization with respect to a large number of design variables requires a gradient-based optimization algorithm and an efficient approach for computing these gradients.

One of the most popular methods for computing derivatives is the finite-difference method. However, the computational cost of this method does not scale well with the number of design variables, and it is subject to large numerical errors due to subtractive cancellation. The complex-step method is more accurate but does not scale any better [51]. When using the adjoint method, however, the cost of computing derivatives in a gradient is practically independent of the number of design variables. For a review on derivative computation methods, see Martins and Hwang [52] or, for a more specialized review focused on CFD, see Peter and Dwight [53].

A function of interest (I) for an aerodynamic shape optimization (which could be an objective or constraint function) is dependent on the design variables describing the wind turbine blade geometry ($\boldsymbol{x}$), and the flow state variables ($\boldsymbol{w}$). Such a function depends both directly on the design variables, and indirectly through the solution of the governing equations. The direct dependence assumes that the state variables remain constant, while the indirect dependence is due to the change in the state variables required to satisfy the governing equations for a given set of design variables. Hence the total derivative of the function of interest can be written as,

$$\frac{dI}{d\boldsymbol{x}} = \frac{\partial I}{\partial \boldsymbol{x}} + \frac{\partial I}{\partial \boldsymbol{w}}\frac{d\boldsymbol{w}}{d\boldsymbol{x}}. \tag{4}$$

The total derivative of the flow solution with respect to the design variables ($d\boldsymbol{w}/d\boldsymbol{x}$) can be found by observing that the total derivative of the CFD residuals $\boldsymbol{R}(\boldsymbol{x}, \boldsymbol{w}(\boldsymbol{x})) = 0$ with respect to $\boldsymbol{x}$ remains zero for a feasible solution, and thus,

$$\frac{d\boldsymbol{R}}{d\boldsymbol{x}} = \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}} + \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{w}}\frac{d\boldsymbol{w}}{d\boldsymbol{x}} = \boldsymbol{0}. \tag{5}$$

We can use the above equation to write a linear system whose solution is the total derivative of the flow solution with respect to the design variables, i.e.,

$$\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{w}}\frac{d\boldsymbol{w}}{d\boldsymbol{x}} = -\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}. \tag{6}$$

Substituting the solution of this linear system into equation (4), we obtain

$$\frac{dI}{d\boldsymbol{x}} = \frac{\partial I}{\partial \boldsymbol{x}} - \frac{\partial I}{\partial \boldsymbol{w}}\left[\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{w}}\right]^{-1}\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}. \tag{7}$$

Noting that the matrix $\partial \boldsymbol{R}/\partial \boldsymbol{w}$ can also be factorized with $\partial I/\partial \boldsymbol{w}$, we can first solve the system of linear equations,

$$\left[\frac{\partial \boldsymbol{R}}{\partial \boldsymbol{w}}\right]^T \boldsymbol{\psi} = -\left[\frac{\partial I}{\partial \boldsymbol{w}}\right]^T. \tag{8}$$

where $\boldsymbol{\psi}$ are called the *adjoint variables*. In this equation, the design variables $\boldsymbol{x}$ do not appear explicitly, and the adjoint vector does not depend on the design variables. This constitutes the major advantage of the adjoint method: the cost of computing the gradient $dI/d\boldsymbol{x}$ is independent of the size of the vector of design variables $\boldsymbol{x}$. After solving for the adjoint variables, we substitute them into the total derivative (4) that we ultimately want to compute, yielding

$$\frac{dI}{d\boldsymbol{x}} = \frac{\partial I}{\partial \boldsymbol{x}} + \boldsymbol{\psi}^T \frac{\partial \boldsymbol{R}}{\partial \boldsymbol{x}}. \tag{9}$$

The equations derived above represent the discrete adjoint approach for computing derivatives, where we are linearizing the discrete system of equations $\boldsymbol{R} = 0$, and then solving the linear equation. This approach contrasts with the continuous approach, where the partial differential equations are linearized first and then discretized [54].

In the adjoint equation (8) and total derivative equation (9), the partial derivatives must be computed. We derive the code to compute these partial derivatives by using the algorithmic differentiation tool Tapenade [55], which automatically parses source code and creates additional code that computes derivatives. Algorithmic differentiation relies on a systematic application of the chain rule to each line in the source code. In this paper, we use the ADflow adjoint implementation of Lyu et al. [46].

2.5 Optimization algorithm

As previously mentioned, a gradient-based optimization algorithm is required to handle the large numbers of shape variables in this shape-optimization problem. The use of such an algorithm, together with the adjoint method, allows us to solve the design optimization problems presented herein.

We use the sparse nonlinear optimizer (SNOPT), which is a sequential quadratic programming (SQP) algorithm capable of solving large-scale nonlinear optimization problems with thousands of constraints and design variables [56]. SNOPT uses a smooth augmented Lagrangian merit function and the Hessian of the Lagrangian is approximated by using a limited-memory quasi-Newton method. The convergence criterion of SNOPT is quantified by the ratio between the largest component of the gradient of the Lagrangian merit function and the L2-norm of the vector of the Lagrange multipliers. The feasibility condition of SNOPT is the ratio of the largest violation of a constraint function with respect to the L2-norm of the vector of design variables [57].

2.6 Aerodynamic shape optimization framework integration

All components just described are integrated into a computational framework capable of optimizing the aerodynamic shape, as shown in Figure 3. ADflow computes a steady-state flow solution for a given shape (corresponding to a specific set of design variables $\boldsymbol{x}$) and also computes aerodynamic performance metrics such as the torque coefficient. The adjoint solver then computes the total derivatives of the quantities of interest with respect to the design variables, as described in Sec. 2.4.
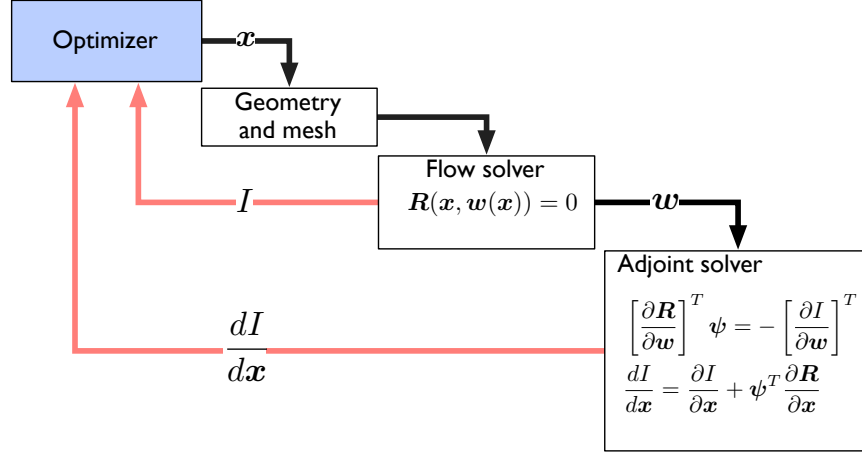

Figure 3: Aerodynamic shape optimization procedure.

Given the values of the objective and constraint functions, as well as the gradients of these functions computed by using the adjoint solver, the optimization algorithm (SNOPT) determines the next set of design variables. The changes in shape are applied to the CFD mesh by using the mesh deformation described in Sec. 2.2, and then a new CFD solution is performed, yielding new values for the objective, constraints, and corresponding gradients. This cycle constitutes a single design optimization iteration and is repeated until both the optimality and feasibility conditions are met.

From the implementation point of view, the components in this framework are integrated by using the Python programming language, which facilitates the coupling of the various codes. The

underlying codes are programmed in Fortran for the best possible computational performance, but each component is wrapped in Python to access the functions required for the framework integration. To setup the optimization, we use pyOpt, which wraps several optimization codes including SNOPT [58]. One critical feature of this coupling is that no reading or writing to file occurs, and all the data transfer between codes is done through memory, making it computationally efficient.

3 Design optimization problem formulation

A sound formulation for the design optimization problem is crucial to obtain realizable designs. The choice of objective function is one of the key decisions in this formulation. Currently, most wind turbine optimization procedures use a systems engineering approach in which the cost of energy is the key factor for optimization.

However, given the large number of disciplines involved in determining the cost of energy, it is not always possible to directly optimize this cost. In this work, the objective is to maximize the energy production below the rated power for a given cost, more specifically the TCC. In our optimizations, we keep the TCC approximately the same as the baseline by keeping the rotor size constant. The baseline turbine is the NREL VI wind turbine, a two-bladed turbine developed to provide validation data [41].

We know that the NREL VI wind turbine performance is not optimal, since the thrust values are rather small for the specified flow conditions. Unlike most modern wind turbines, the NREL VI wind turbine blade rotates at a constant angular velocity for every free-stream velocity. Therefore, the generated power is only dependent on torque [59], so maximizing the torque directly maximizes the generated power. Thus, the objective function of our design optimization is the torque coefficient C_q, which we seek to maximize.

In this work, we use three types of design variables: shape, pitch, and twist. The shape variables s apply perturbations to the surface of the blade in the x direction (see Figure 1), thus controlling the airfoil shape, and allowing the airfoil shapes to change in the spanwise direction. The twist variables ξ twist the airfoils about the 25% chord line and are distributed along the spanwise direction. The pitch variable θ twists the whole blade as a solid body and is equivalent to adding a constant to all twist variables.

The twist and pitch are potentially redundant, because pitch can be expressed as a linear combination of twist variations. This redundancy is problematic for optimization because it allows for an infinite number of valid solutions. To avoid this problem, the twist variable closest to the hub is fixed, and the twist distribution is expressed relative to this fixed value. The other alternative would be to remove the pitch variable and include twist closest to the hub. However, using pitch is more intuitive to wind turbine designers and, in addition, explicitly including pitch results in superior optimization. Finally, we do not include blade radius or chord variables in our study because varying these variables impacts the TCC in ways that we do not model.

Given that the turbine blade shape optimization does not consider any structural sizing, we need to add constraints to avoid results that are physically unrealizable. To this end, we impose thickness constraints on the blade to ensure that there is adequate depth for the structural box. The location of the structural box is assumed to be between 15% and 50% of the local blade chord. We impose constraints in the structural-box region such that the airfoil profiles cannot be thinner than the original shape.

In this work, we solve a sequence of design optimization problems (whose results are presented in Sec. 5) developed to progressively gain an understanding of the design optimization problem and to gradually verify our approach. One of the features we change between each problem is the

set of design variables. We use four distinct sets: The first set consists of just the pitch angle. In the second problem, we solve a design optimization problem with twelve twist variables distributed along the span of the blade and, in the third problem, we optimize only airfoil shapes when using 240 shape variables. The fourth set combines the twist and shape variables and constitutes what we consider to be the most complete set.

In addition to changing the design variables, we also optimize for a single wind speed, and then progress to considering multiple wind speeds. The single-speed optimization consists in maximizing the torque coefficient C_q at 7 m/s. The optimization problem is stated mathematically as follows:

$$\max_{\theta, \boldsymbol{\xi}, \boldsymbol{s}} \quad C_q\left(\theta, \boldsymbol{\xi}, \boldsymbol{s}\right)$$
$$\text{subject to} \quad 1 \leq \bar{\boldsymbol{t}} \leq 3, \tag{10}$$

where $\bar{\boldsymbol{t}}$ is the vector of local thicknesses normalized with respect to the corresponding local baseline thickness values.

Because wind turbines must perform well for a range of wind speeds, multiple-speed optimization is also performed to maximize the power generation for a range of wind speeds. The multiple-speed optimization maximizes a weighted average of the torque coefficients at wind speeds from cut-in to rated and take the values $U_i = [5, 6, 7, 8, 9]$ m/s. The multiple-speed optimization can be written as

$$\max_{\theta, \boldsymbol{\xi}, \boldsymbol{s}} \quad \sum_{i=1}^{5} w_i C_q\left(\theta, \boldsymbol{\xi}, \boldsymbol{s}, U_i\right)$$
$$\text{subject to} \quad 1 \leq \bar{\boldsymbol{t}} \leq 3. \tag{11}$$

Although the weights w_i in the weighted average are all the same here, they could be changed according to the frequency distribution of wind speeds at a certain location to yield a more tailored design [7].

4 Verification and validation

To have confidence in our design optimization results, we verify the convergence of our RANS solver and validate it by using the experimental results of the NREL VI turbine blade. Our RANS solver had difficulty converging when doing a CFD analysis with the original NREL VI geometry. We determined that this was due to the cylindrical shape at the blade root, which causes massively separated flow, and thus the steady RANS solver did not converge to the desired tolerance. To address this issue, we removed the cylindrical section at the blade root that starts from the hub center and extends 1.257 m in the spanwise direction. Thus, we only modeled the portion of the blade with an airfoil shape. With the cylindrical portion of the blade removed, steady CFD solutions converge well. Removal of the cylindrical section results in a small decrease in the computed torque (by 2%–5% for fully attached flow conditions, depending on wind speed).

We use three meshes with an increasing level of refinement both to make the most of the available computational resources and to verify the results via CFD. These meshes are listed in Table 1. The level 2 (L2) mesh is the coarsest and is used in our first optimization runs to make sure that the problem is well posed. Table 1 also lists the torque coefficient C_q and thrust coefficient C_t, which are computed with these meshes. The torque coefficient computed with the coarse mesh is 29% greater compared with the torque coefficient computed with the finest mesh, whereas the fine mesh differs by 8.7%. Although this difference seems significant, as shown in previous work,

the improvement in the performance is accurately quantified even when the performance is not accurately quantified [35].

Level	Mesh cells	C_q	C_t
L0 (finest)	22×10^6	0.4258	0.6529
L1 (fine)	2.6×10^6	0.3889	0.6667
L2 (coarse)	325×10^3	0.3019	0.5791

Table 1: Mesh convergence of the modified NREL VI blade for 7 m/s wind.

To validate our CFD analysis against the experimental data provided by NREL [60], we compute the torque for wind speeds ranging from cut-in to cut-out, and plot both sets of results in Figure 4. For fully attached flow—in wind speeds ranging from 5 to 7 m/s—our numerical results agree with the experimental results within the experimental error. At 10 m/s the flow is separated, and the steady CFD does not yield reliable results. This is expected, since steady RANS becomes less accurate as the separated region increases. However, most optimizations presented in Sec. 5 are for a wind speed of 7 m/s. The exception is the multiple wind speed optimization, where we include two speeds above 7 m/s. We include the torque computations both with and without the root cylindrical section to show that removing this section does not significantly affect the accuracy.
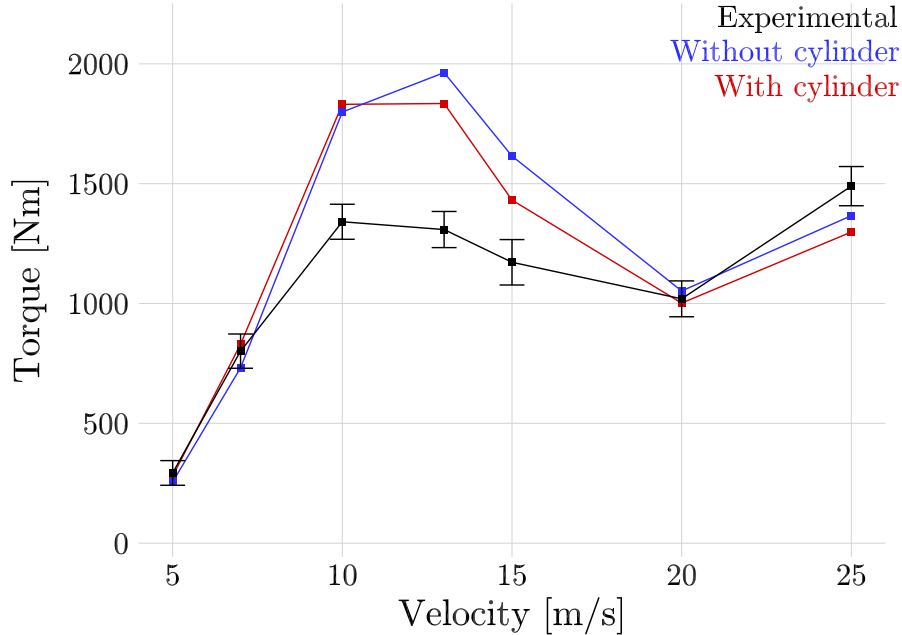

Figure 4: The computed torque values compare well against experimental data until the onset of separation, and the impact of removing the cylindrical root section is not significant.

5 Results

We now present the solution of the optimization problems described in the previous section and discuss the results. We present a series of designs with increasing complexity and optimize the blade design with respect to blade pitch angle, twist angle, and airfoil shape variables for single and multiple wind speeds. We compare all the results to obtain an intuitive understanding of blade design optimization. Unless otherwise stated, the results are obtained with the L2 (coarse) mesh.

5.1 Single-wind-speed optimization

This section presents the optimization results obtained by solving the optimization problem (10)at a fixed wind speed of 7 m/s. We solve a sequence of three problems with an increasing number of design variables, culminating in a case that includes the complete set of variables.

5.1.1 Pitch angle

For a first optimization, we only optimize the blade pitch angle. As shown in Figure 5, only five optimization iterations are required to satisfy the optimality conditions. The optimization increases the pitch angle by 5.14 degrees, which yields an increase of 5.1% in torque.
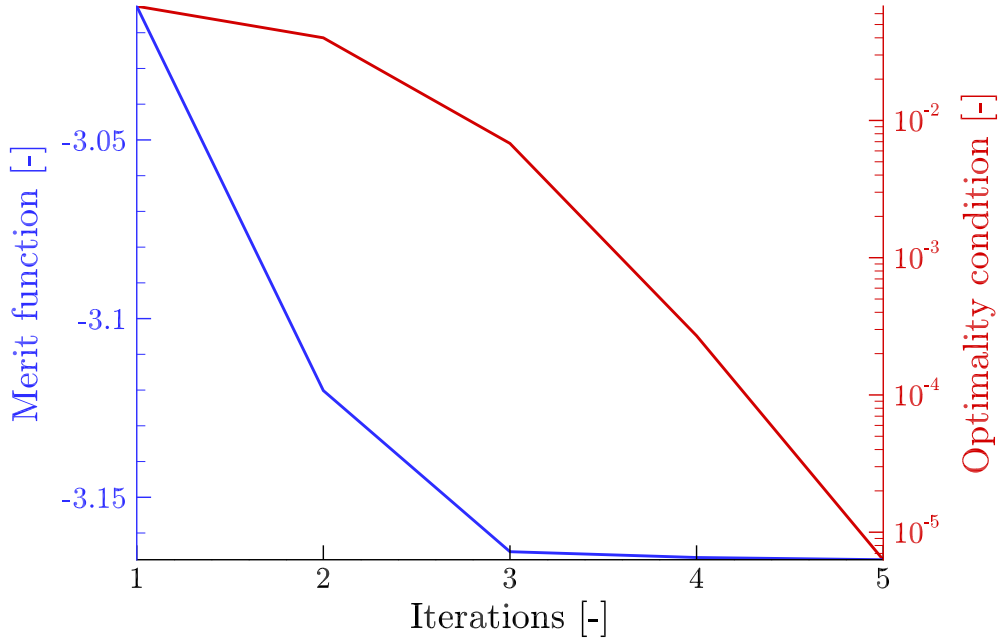

Figure 5: Optimization history with pitch angle as design variable.

Figure 6 shows that the magnitude of the optimized pressure distributions increases, resulting in a higher torque. The largest increase in torque is near the tip, which is expected because most of the power extraction occurs near the tip. The sensitivity contours shown in Figure 7 are consistent with this result, because the largest sensitivity occurs in the tip region.

The largest pressure difference is at the upper part of the leading edge of the blade, whereas the pressure coefficient at the trailing edge does not change significantly compared with the baseline design. The magnitude of the negative torque at the tip of the original blade, which is due to

tip losses, increases during optimization, but this has a positive overall impact on maximizing the torque by using a single pitch value for the entire blade.
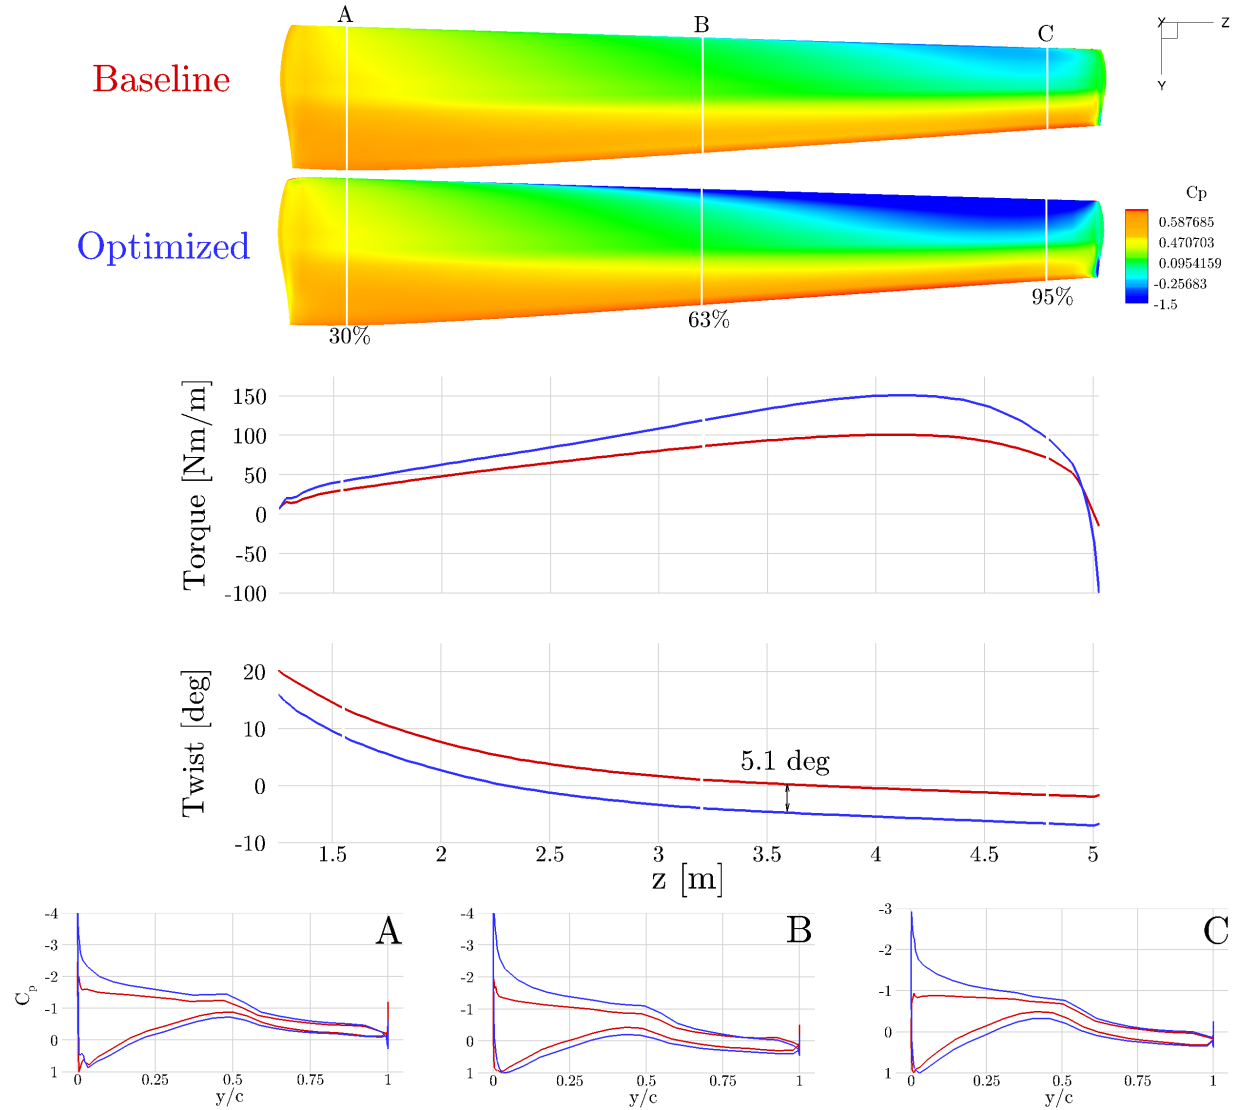

Figure 6: NREL VI baseline blade (red) compared to the blade optimized with respect to pitch angle (blue).

5.1.2 Pitch angle and twist variables

The second optimization study adds the spanwise twist distribution to the blade pitch angle, and optimization is done with 11 sections relative to pitch angle. The resulting twist distribution is shown in Figure 8, which shows that the torque coefficient increases by 5.9% compared with the baseline design. This corresponds to less than 1% improvement in performance compared with the pitch-only optimization. Note that the twist angle reduces abruptly in the tip region, although this localized feature has no significant effect overall on the optimized results.

As seen from Figure 8, the optimized twist distribution oscillates. This result is typical of cases where too few design variables are used. As we shall see in the next section, this problem vanishes
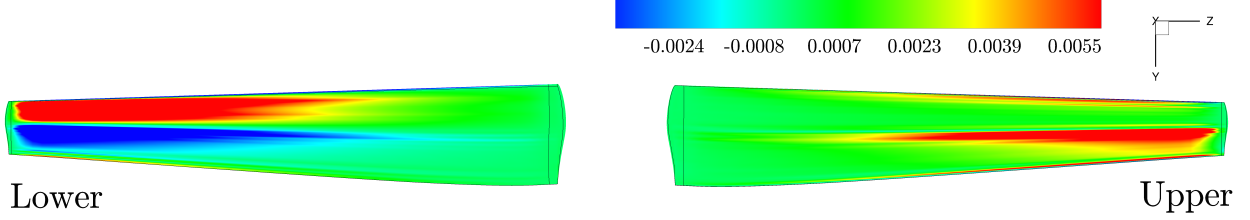
13

Figure 7: Torque coefficient sensitivity ($\partial C_q / \partial \boldsymbol{s}$) distribution on lower and upper surface of NREL VI blade.

when airfoil shape variables are added.

5.1.3 Pitch angle, twist, and shape variables

In conventional wind turbine blade design methods, designers choose from a limited set of predefined airfoil shapes. CFD-based aerodynamic shape optimization enables the design of fully customized airfoil shapes along the span.

For this design optimization, we added airfoil shape variables to the 11 twist variables and the blade pitch angle used in the previous optimization. The airfoil shape is controlled at the 12 spanwise stations described in Sec. 2.1. The total number of design variables (shape, twist, and pitch) is thus 252.

The results of this design optimization are shown in Figure 9. The torque coefficient increases by 31%, which is significantly greater than the 5.9% increase obtained in the previous optimization. Additionally, the twist distribution is smoother, but the abrupt change in twist at the tip persists because of the relatively coarse CFD mesh.

To resolve this problem of abrupt change in twist at the tip and thereby obtain a better overall result, we repeat the same design optimization but with the L1 (fine) CFD mesh. The resulting optimum is shown in Figure 10, which shows that the spanwise oscillation present in the results obtained with the coarse mesh completely disappears when using the fine mesh. In addition, the twist no longer changes abruptly at the tip. These results confirm that the L1 mesh sufficiently resolves the complexity in the flow field. As a result, the distribution of the pressure coefficient is now smooth, which is consistent with a well-resolved flow solution.

When using the L1 mesh, the torque coefficient increases 22.4% relative to the baseline torque of the NREL VI blade. This increase in torque is less than the 31% obtained with the coarse mesh, which reinforces the point that optimization must be done with a sufficiently fine mesh.

As shown in Figure 10, the airfoils optimized with this fine mesh exhibit more camber, and the root region has a highly cambered trailing edge. The optimizer converged to this shape to increase the lift in the absence of high wind speeds near the root. Figure 11 shows the thrust of the optimized design with 69% increase in the outboard region of the blade. This also increases the blade flapwise bending moment, so the blade structure should be reinforced were this blade constructed.

For future design optimization studies, we recommend considering how aerodynamic shape optimization affects structural design. This could be achieved by enforcing a bending moment constraint to limit blade mass. A better approach, however, is to consider the trade-offs between structural weight and aerodynamic performance. This requires modeling the blade aeroelasticity, which has been done for aircraft wing design [44, 61], and doing the same for wind turbine blade
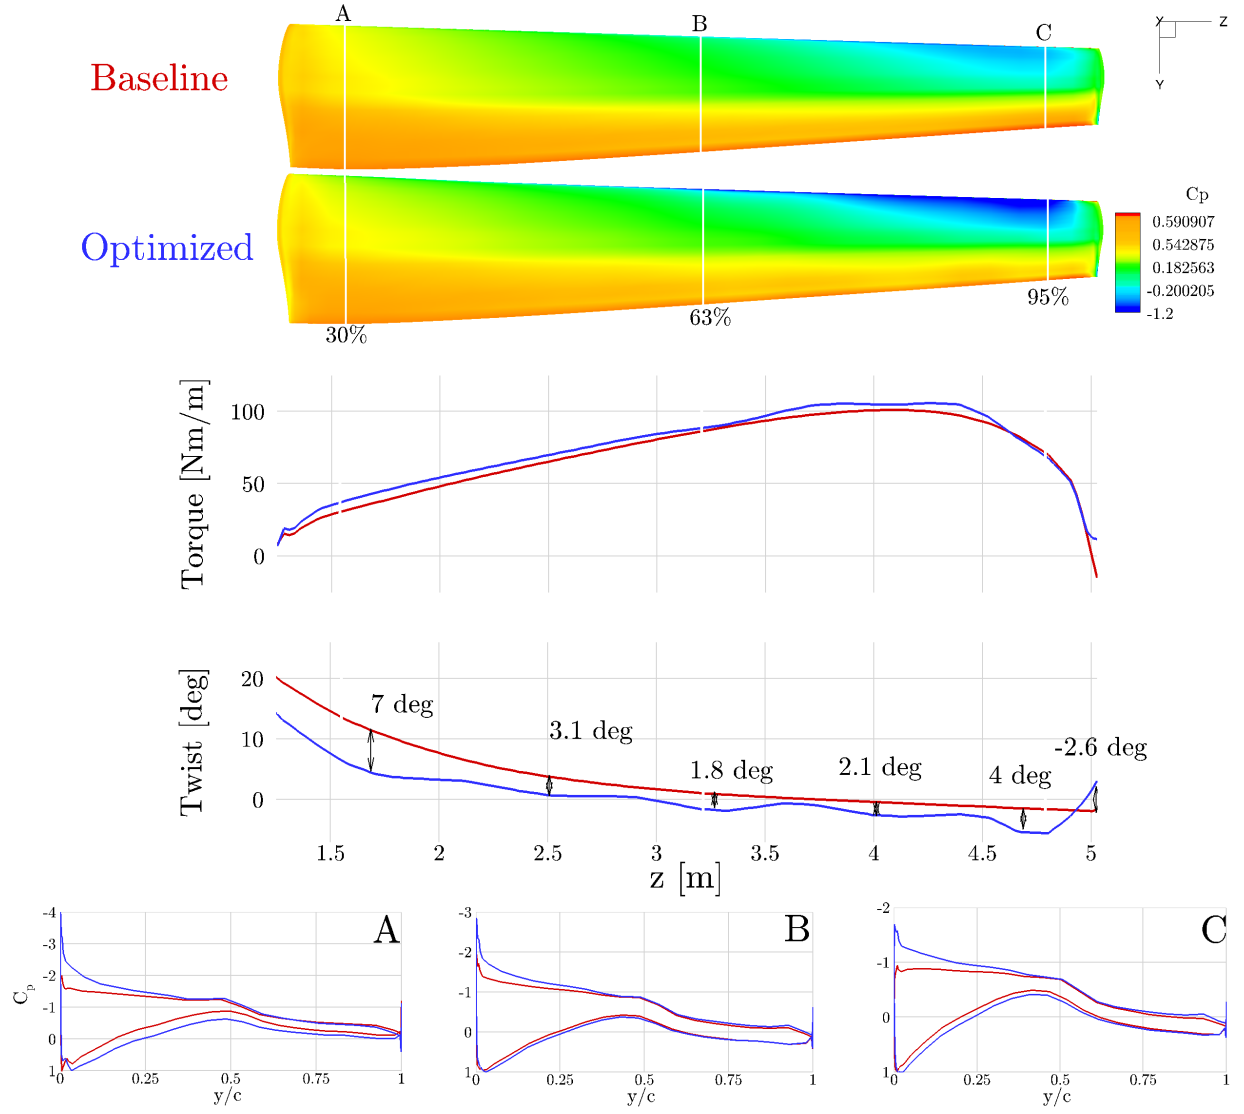
14

Figure 8: NREL VI baseline blade (red) compared to the twist optimized blade using the coarse (L2) mesh (blue).
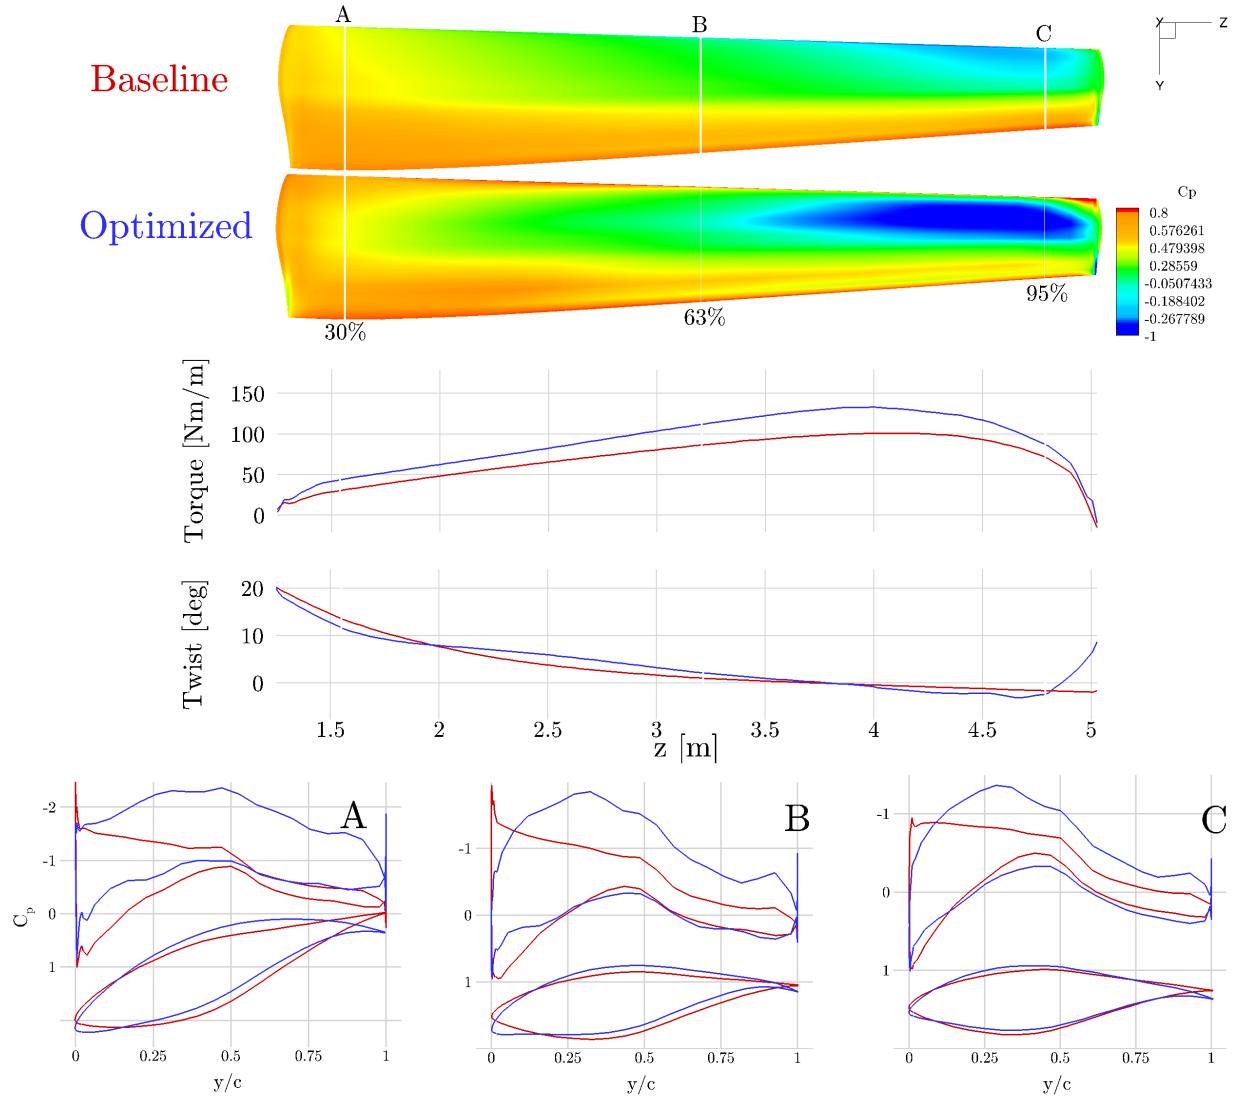
Figure 9: NREL VI baseline blade (red) compared to the blade optimized with respect to pitch angle, twist, and airfoil shape variables using the coarse (L2) mesh (blue); showing chordwise C_p distributions and airfoil shapes for cross sections A–C.
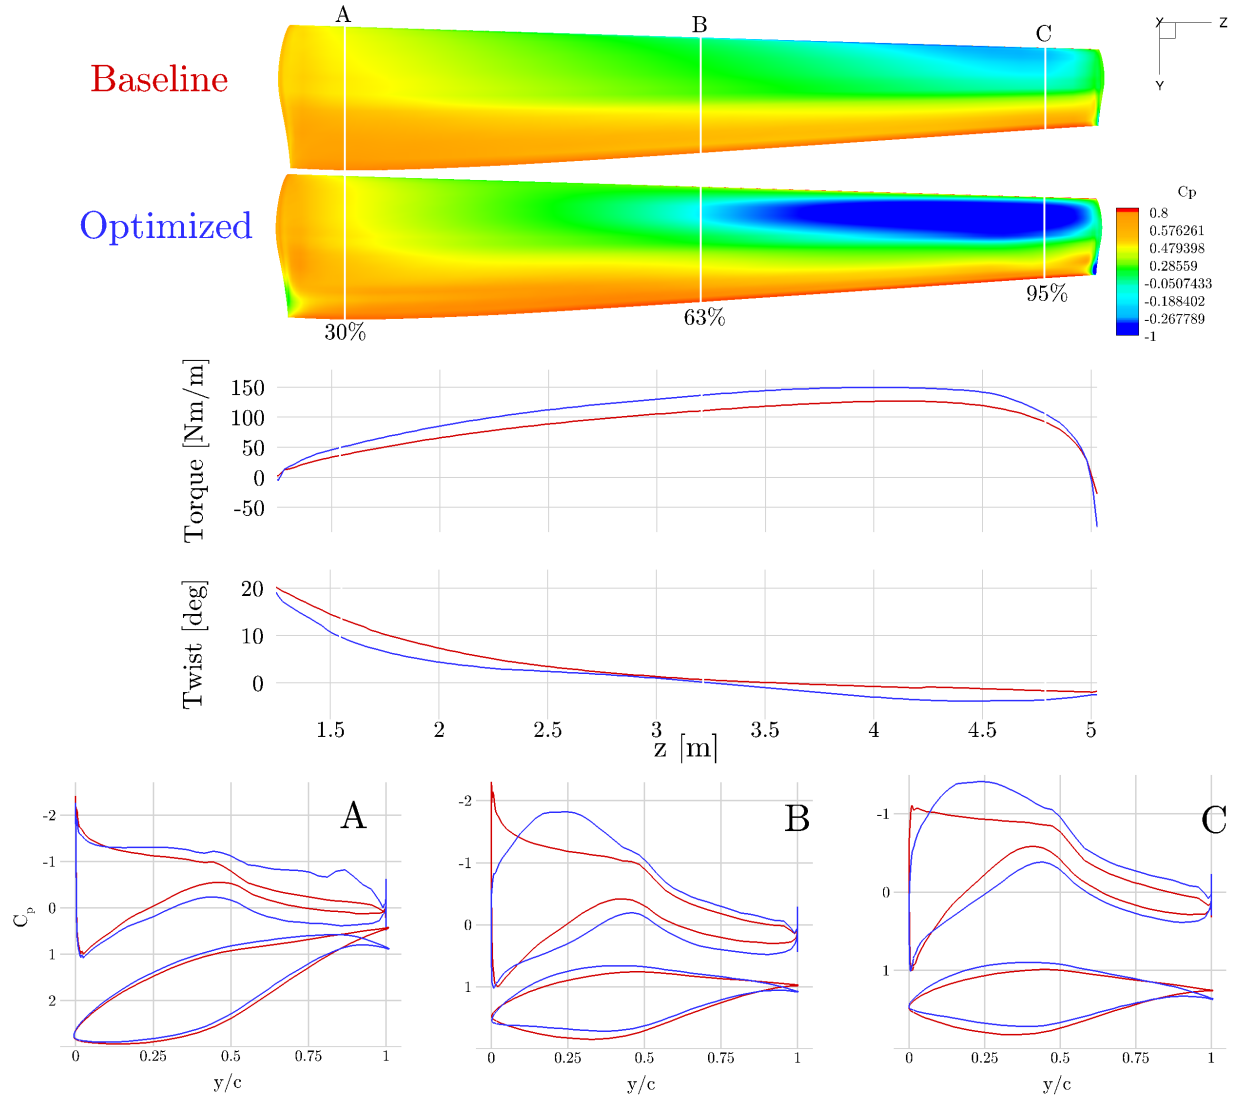
Figure 10: NREL VI baseline blade (red) compared to the blade optimized with respect to pitch angle, twist, and airfoil shape variables using the fine (L1) mesh (blue); showing chordwise C_p distributions and airfoil shapes for cross sections A–C.
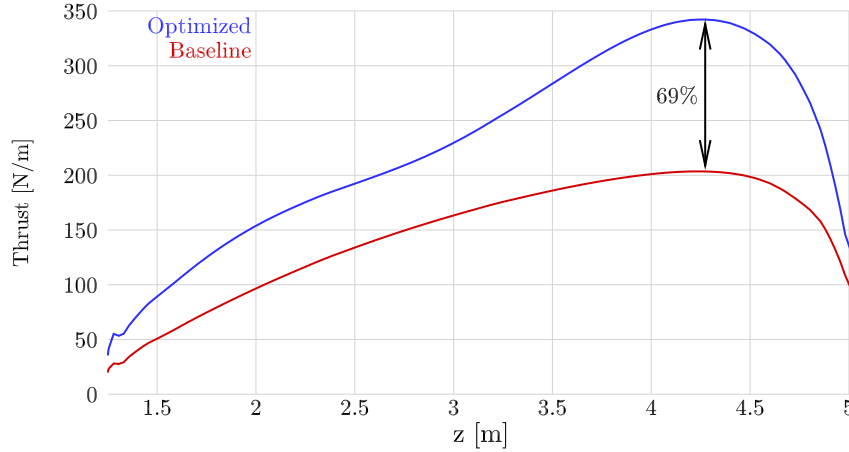
design should be possible.

Figure 11: The optimized blade thrust spanwise distribution (blue) exhibits a large thrust increase compared to the baseline (red).

The L0 mesh results in highly accurate CFD simulation results that are within the standard deviation of the experimental results (96% accurate with respect to averaged torque value). However, the computational cost incurred when using this mesh is too high for performing optimization. Therefore, we used the L0 mesh only for analysis, that is, we analyzed the optimum shape obtained with the L1 mesh using the L0 mesh. Given the increase in predicted torque for larger mesh resolutions, the increase in optimized torque should be even higher when using the L0 mesh.

We show the result of this analysis in Figure 12 in the form of airfoil pressure distributions only, since the twist and airfoil shapes remain the same. It is clear that the pressure coefficient distributions are even smoother than for the L1 mesh. Also, the pressure distribution at root region is better resolved, and the root vortex can be observed. The L0 mesh analysis estimates a 29.1% increase in torque, which is 6.7% more than the increase predicted with the L1 mesh. This can be explained by the fact that the tip and root vortices are better resolved, and that the torque is predicted more accurately when using the finest mesh.

5.2 Optimization for multiple wind speeds

The design optimization results reported in the previous sections consider only a single wind speed. However, real wind turbines must operate over a wide range of wind speeds that may range anywhere between cut-in to cut-out. Therefore, the design should be optimized for multiple wind speeds and by using site-specific wind profiles [62, 7].

We solve the optimization problem (11) with five different wind speeds, ranging from cut-in (5 m/s) to the rated wind speed (9 m/s), with 1 m/s increments (region-2 operation). The weights associated with each wind speed are equal, so the objective function is the average torque coefficients for these wind speeds. Given the propagation of flow separation from the root to tip of the NREL VI blade at 5 to 10 m/s due to rotational effects, it is expected that the accuracy of the flow solution decreases with higher wind speeds. As a result, it is assumed that this multiple speed optimization yields valid results for wind speeds between 5 and 9 m/s. The optimization is done with the L1 (fine) mesh and the blade pitch, twist, and shape variables.
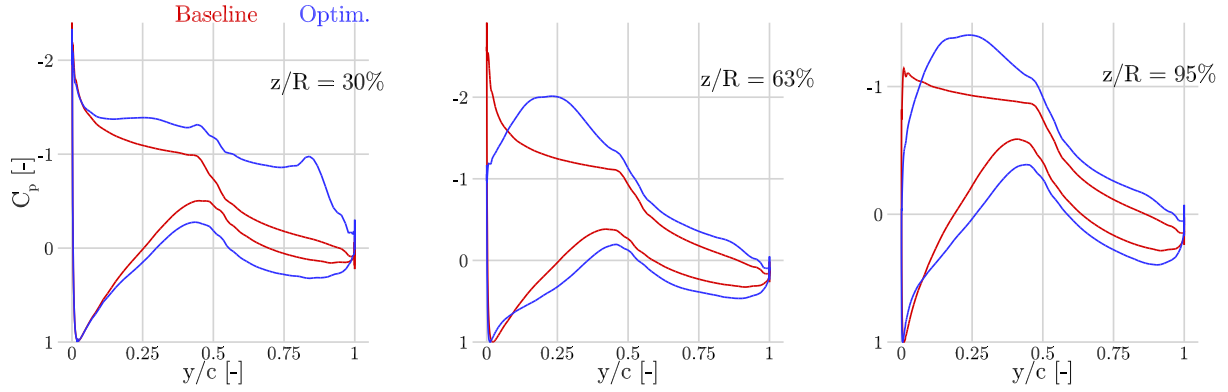
18

Figure 12: Chordwise C_p distributions for the fine mesh (L1) optimized geometry analyzed using the finest (L0) mesh (blue), compared to the baseline case analyzed using the same level mesh (red).

The multiple-speed-optimized blade is shown in Figure 13. The multiple speed design exhibits characteristics similar to those of the single-speed design. The camber is increased over the entire span of the airfoils, with the airfoil near the root exhibiting extreme cambering at the trailing edge.

Figure 14 shows the results for a wind speed of 7 m/s for single- and multiple-speed torque, twist, and distribution of airfoil pressure at 30%, 63%, and 95% of the blade length. The multiple-speed optimization converges to a lower twist in the mid sections of the blade, so the torque is slightly less than for the single-speed design. The pressure coefficient distributions and the airfoil shape are approximately the same, with a small increase in the camber of the multiple-speed airfoils.

Figure 15 compares the torque obtained from multiple-speed-optimized design with that obtained from the baseline design. The torque averaged over all wind speeds increases by 22.1% compared with that of the baseline design. Using these torques to compute power output at a fixed angular velocity of 7.5 rad/s allows us to calculate the AEP assuming a Rayleigh probability density function with a shape factor of 2 and a scale factor of 9.47[1]. The optimized NREL VI blade delivers a 22.2% increase in AEP relative to the baseline design below the rated wind speed.

5.3   Computational performance

One of the main motivations of this work was to make RANS-based aerodynamic shape optimization of wind turbines a tractable method with fast turnaround by using high-performance computing resources. Table 2 summarizes the computational performance of the optimization studies. These cases were run in parallel on 256 processors.

All CFD simulations converge by six orders of magnitude with respect to the residual of density. The optimality tolerance achieved by the SNOPT optimizer varies, as shown in Table 2; this variation in tolerance is due to noise in the torque coefficient, whose fourth significant digit fluctuates.

These optimizations take between less than two hours of computation time (for single-speed optimization using the coarse mesh) and just over one day (for multiple-speed optimization using the fine mesh). These times are short enough for designers to use this approach in practice.
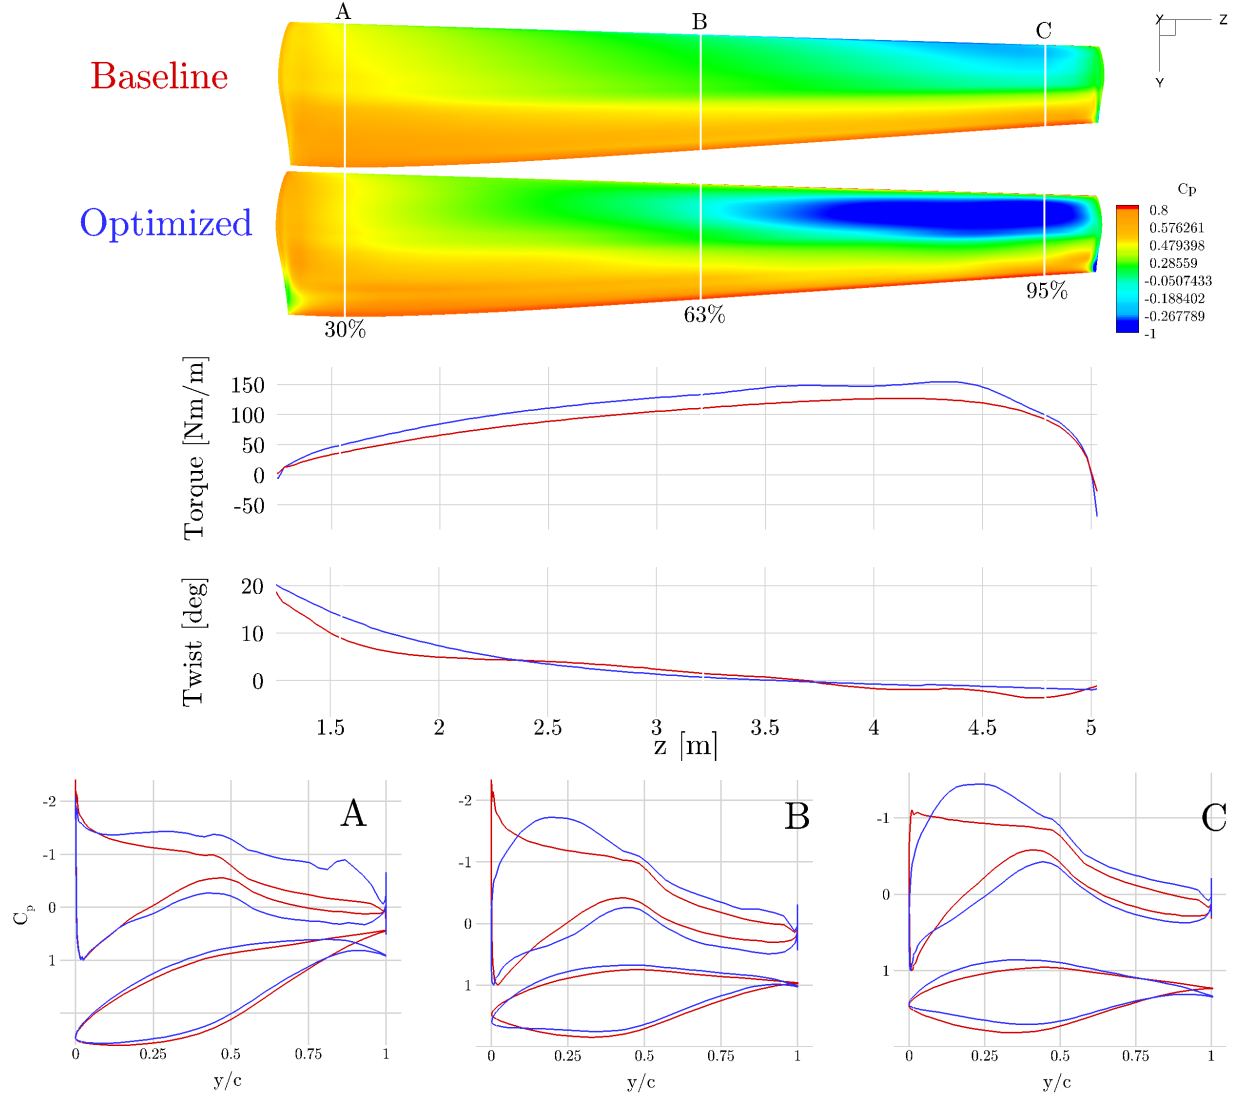
Figure 13: Multiple speed optimization result analyzed at a wind speed of 7 m/s. Red line is for the baseline case, and blue line is for the optimized case. For cross sections A–C, the upper figure shows the C_p, and the lower figure shows the airfoil cross section.

Table 2: Computational performance of aerodynamic shape optimization run on 256 computing cores.

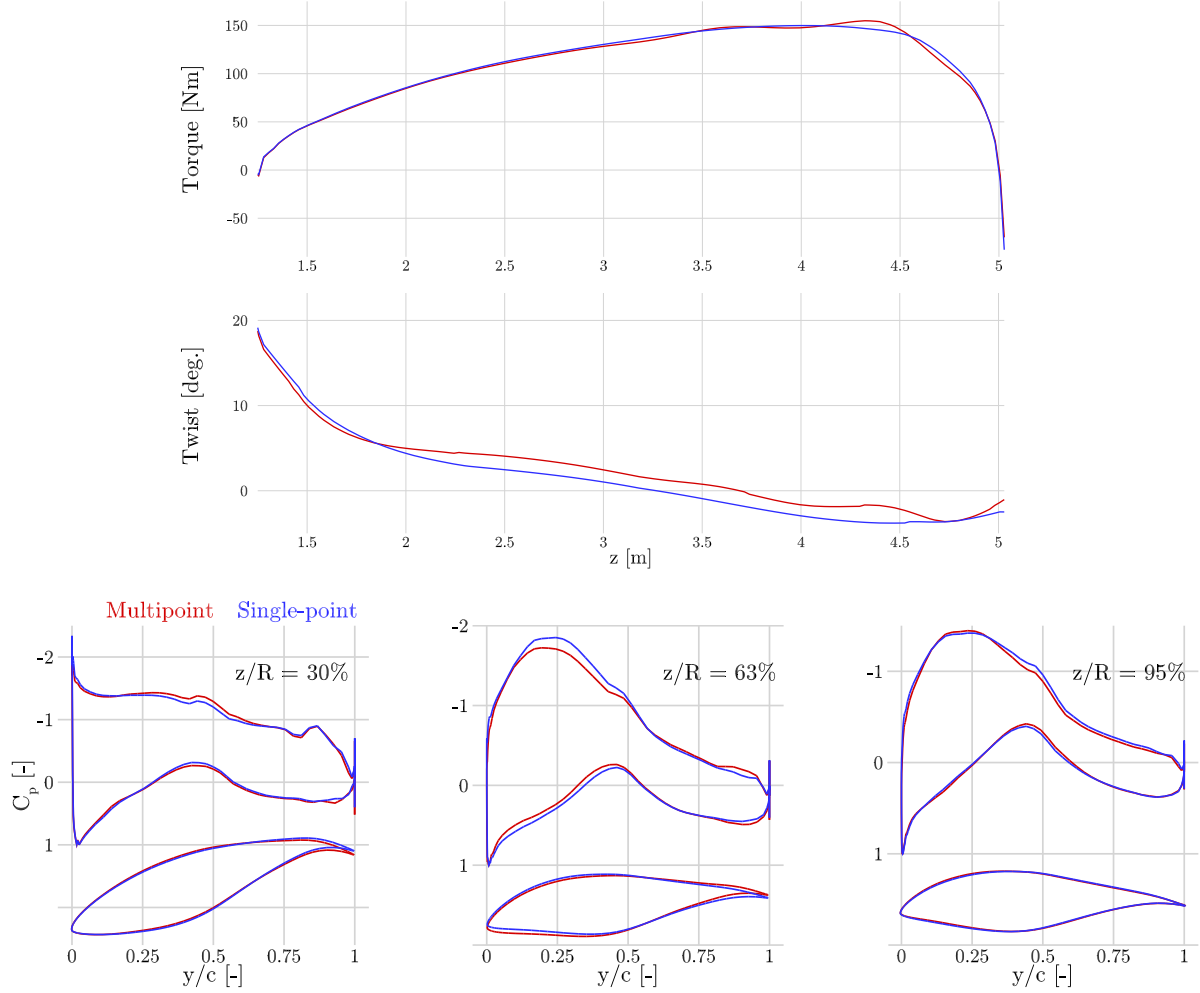| Optimization case	Iterations	Optimality	Wall time [hr]
Single-speed pitch, twist, and shape (L2)	23	$10^{-1.89}$	1.58
Single-speed pitch, twist, and shape (L1)	9	$10^{-1.94}$	8.25
Multiple-speed pitch, twist, and shape (L1)	15	$10^{-1.42}$	27.75

Figure 14: Comparison of multiple-speed (red) and single-speed (blue) optimized designs at 7 m/s.

6 Conclusion

This work shows that RANS-based aerodynamic shape optimization of wind turbine blades that considers the detailed shape of the entire blade is now possible and results in significant improvements in performance. To demonstrate the capability of our framework, we started with the NREL VI wind turbine blade as a baseline and maximize the torque coefficient with respect to blade shape. Thickness constraints prevent the blade from getting any thinner than the NREL VI blade. The optimization results in a 22.4% increase in the torque coefficient.

   With computational times ranging from under two hours up to just over one day when using 256 processors, such optimizations are now sufficiently efficient to be used routinely by wind turbine designers. The efficiency of this approach to optimize the aerodynamic shape hinges on a fast CFD solver and the use of gradient-based optimization in conjunction with a method to accurately compute gradients efficiently by using an adjoint method. Robust geometry parametrization and mesh deformation are also critical.
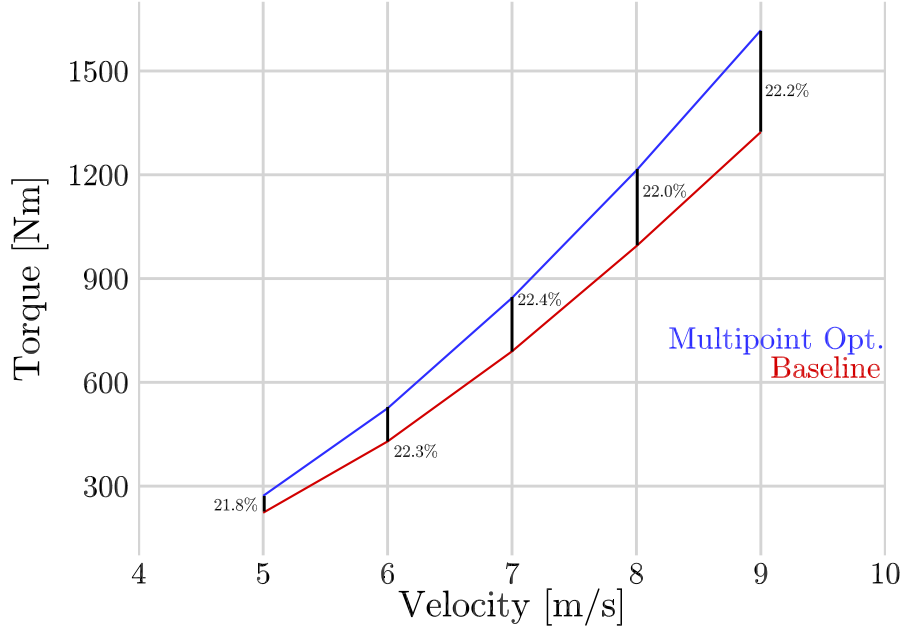
21

Figure 15: Torque as a function of wind velocity for baseline (red) and multiple-speed optimized blade (blue).

7 Future work

While we have shown that it is possible to perform aerodynamic shape optimization based on steady-state RANS, unsteady effects due to tower shadow and massive flow separation are present and should be modeled to obtain truly high-fidelity results. To overcome this limitation, we could model the separated unsteady flow at the higher speeds using LES, detached eddy simulations, or unsteady RANS. Given the high computational cost of the flow solutions using these approaches, together with the cost of solving the unsteady adjoint equations, it is currently not viable to include them in an optimization iteration.

Variable-speed operations of wind turbines with a controller could also be performed. This can be implemented by adding the partial derivatives in adjoint equations corresponding to the rotational speed. Constraints on the rotational speed variable can then be implemented as well. For example, a constant rotational speed can be imposed at a certain wind speeds to prevent structural failure.

Other avenues of developments would be to do static aeroelastic (aerostructural) optimization, as already done for aircraft-wing design [44, 61]. Such an approach would eliminate the need for explicit thickness constraints and take into account the trade-off between aerodynamic performance and weight, resulting in optimal passive load alleviation.

The high-fidelity aerodynamic shape optimization developed in this research can be extended to other rotational problems such as helicopters, gas turbines, and propellers for both aircraft and marine applications.

Acknowledgments

References

[1] Ashuri, T. and Zaaijer, M. B., "Size effect on wind turbine blades design drivers," *The proceedings of the European Wind Energy Conference and exhibition, Brussels, Belgium*, European Wind Energy Association, 2008, pp. 1–10.

[2] Ashuri, T. and Zaaijer, M. B., "The future of offshore wind turbines," *3rd PhD seminar on wind energy in Europe, Pamplona, Spain*, European Academy of Wind Energy, 2007, pp. 1–7.

[3] Ning, A., Damiani, R., and Moriarty, P., "Objectives and Constraints for Wind Turbine Optimization," *Journal of Solar Energy Engineering*, Vol. 136, No. 4, Nov. 2014, pp. 041010–1–12. doi:10.1115/1.4027693.

[4] Martins, J. R. R. A. and Lambe, A. B., "Multidisciplinary Design Optimization: A Survey of Architectures," *AIAA Journal*, Vol. 51, No. 9, September 2013, pp. 2049–2075. doi:10.2514/1.J051895.

[5] Ashuri, T., *Beyond Classical Upscaling: Integrated Aeroservoelastic Design and Optimization of Large Offshore Wind Turbines*, Ph.D. thesis, Delft University of Technology, the Netherlands, 2012.

[6] Fuglsang, P. and Madsen, H., "Optimization Method for Wind Turbine Rotors," *Journal of Wind Engineering and Industrial Aerodynamics*, Vol. 80, 1999, pp. 191–206.

[7] Kenway, G. K. W. and Martins, J. R. R. A., "Aerostructural Shape Optimization of Wind Turbine Blades Considering Site-Specific Winds," *Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Victoria, BC, Sept. 2008. doi:10.2514/6.2008-6025, AIAA 2008-6025.

[8] Xudong, W., Shen, W. Z., Zhu, W. J., Sørensen, J. N., and Jin, C., "Shape optimization of wind turbine blades," *Wind Energy*, Vol. 12, No. 8, 2009, pp. 781–803. doi:10.1002/we.335.

[9] Ning, A., "A Simple Solution Method for the Blade Element Momentum Equations with Guaranteed Convergence," *Wind Energy*, Vol. 17, No. 9, sep 2014, pp. 1327–1345. doi:10.1002/we.1636.

[10] Ashuri, T., Van Bussel, G., Zaayer, M., and Van Kuik, G., "An analytical model to extract wind turbine blade structural properties for optimization and up-scaling studies," *The Science of making Torque from Wind, Heraklion, Crete, Greece*, 2010, pp. 1–7.

[11] Muskulus, M. and Schafhirt, S., "Design optimization of wind turbine support structures—a review," *Journal of Ocean and Wind Energy*, Vol. 1, No. 1, 2014, pp. 12–22.

[12] Haghi, R., Ashuri, T., van der Valk, P. L. C., and Molenaar, D. P., "Integrated multidisciplinary constrained optimization of offshore support structures," *Journal of Physics: Conference Series*, Vol. 555, No. 1, 2014, pp. 1–10. doi:10.1088/1742-6596/555/1/012046.

[13] Zwick, D., Muskulus, M., and Moe, G., "Iterative optimization approach for the design of full-height lattice towers for offshore wind turbines," *Energy Procedia*, Vol. 24, 2012, pp. 297–304.

[14] Dai, J. C., Hu, Y. P., Liu, D. S., and Long, X., "Aerodynamic loads calculation and analysis for large scale wind turbine based on combining BEM modified theory with dynamic stall model," *Renewable Energy*, Vol. 36, No. 3, 2011, pp. 1095–1104.

[15] Robison, D. J., Coton, F. N., Galbraith, R. A. M. C. D., and Vezza, M., "Application of a prescribed wake aerodynamic prediction scheme to horizontal axis wind turbines in axial flow," *Wind Engineering*, Vol. 19, No. 1, 1995, pp. 41–51.

[16] Ashuri, T., Zhang, T., Qian, D., and Rotea, M., "Uncertainty Quantification of the Levelized Cost of Energy for the 20MW Research Wind Turbine Model," *AIAA Science and Technology Forum and Exposition (SciTech), Wind Energy Symposium, San Diego, California*, American Institute of Aeronautics and Astronautics., 2016, pp. 1–11.

[17] Castillo Capponi, P., Ashuri, T., Van Bussel, G., and Kallesoe, B., "A Non-Linear Upscaling Approach for Wind Turbines Blades Based on Stresses," *European Wind Energy Conference, Brussels, Belgium*, European Wind Energy Association, 2011, pp. 1–8.

[18] Afjeh, A. A. and Keith, T. G., "A vortex lifting line method for the analysis of horizontal axis wind turbines," *Journal of solar energy engineering*, Vol. 108, No. 4, 1986, pp. 303–309.

[19] Ashuri, T., Martins, J. R., Zaaijer, M. B., van Kuik, G. A., and van Bussel, G. J., "Aeroservoelastic Design Definition of a 20 MW Common Research Wind Turbine Model," *Wind Energy*, Vol. 19, No. 11, 2016. doi:10.1002/we.1970.

[20] Ning, A. and Petch, D., "Integrated Design of Downwind Land-based Wind Turbines using Analytic Gradients," *Wind Energy*, Feb. 2016, pp. 1–17. doi:10.1002/we.1972.

[21] Sørensen, N. N. and Hansen, M. O. L., "Rotor performance predictions using a Navier-Stokes method," *American Institute of Aeronautics & Astronautics*, Vol. 98, 1998, pp. 52–59.

[22] Duque, E. P. N., Burklund, M. D., and Johnson, W., "Navier-Stokes and comprehensive analysis performance predictions of the NREL Phase VI experiment," *American Institute of Aeronautics & Astronautics*, Vol. 125, No. 4, 2003, pp. 457–467.

[23] Zahle, F., Sørensen, N. N., and Johansen, J., "Wind turbine rotor-tower interaction using an incompressible overset grid method," *Wind Energy*, Vol. 12, No. 6, 2009, pp. 594–619.

[24] Potsdam, M. A. and Mavriplis, D. J., "Unstructured mesh CFD aerodynamic analysis of the NREL Phase VI rotor," *AIAA paper*, Vol. 1221, 2009, pp. 1–18.

[25] Sørensen, N. N., Bechmann, A., and Zahle, F., "3D CFD computations of transitional flows using DES and a correlation based transition model," *Wind Energy*, Vol. 14, No. 1, 2011, pp. 77–90.

[26] Sezer-Uzol, N., Gupta, A., and Long, L. N., *3-D Time-Accurate Inviscid and Viscous CFD Simulations of Wind Turbine Rotor Flow Fields*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 457–464. doi:10.1007/978-3-540-92744-0_57.

[27] Ashuri, T., Rotea, M., Xiao, Y., Li, Y., and Ponnurangam, C. V., "Wind Turbine Performance Decline and its Mitigation via Extremum Seeking Controls," *AIAA Science and Technology Forum and Exposition (SciTech), Wind Energy Symposium, San Diego, California*, American Institute of Aeronautics and Astronautics, 2016, pp. 1–11.

[28] Ribeiro, A. F. P., Awruch, A. M., and Gomes, H. M., "An airfoil optimization technique for wind turbines," *Applied Mathematical Modelling*, Vol. 36, No. 10, 2012, pp. 4898–4907.

[29] Kwon, H. I., You, J. Y., and Kwon, O. J., "Enhancement of wind turbine aerodynamic performance by a numerical optimization technique," *Journal of mechanical science and technology*, Vol. 26, No. 2, 2012, pp. 455–462.

[30] Elfarra, M. A., Sezer-Uzol, N., and Akmandor, I. S., "NREL VI rotor blade: numerical investigation and winglet design and optimization using CFD," *Wind Energy*, Vol. 17, 2014, pp. 605–626. doi:10.1002/we.1593.

[31] Economon, T. D., Palacios, F., and Alonso, J. J., "A viscous continuous adjoint approach for the design of rotating engineering applications," *AIAA Paper*, Vol. 2580, 2013, pp. 1–19.

[32] Hicks, R. M. and Henne, P. A., "Wing Design by Numerical Optimization," *Journal of Aircraft*, Vol. 15, 1978, pp. 407–412.

[33] Jameson, A., "Aerodynamic Design via Control Theory," *Journal of Scientific Computing*, Vol. 3, No. 3, Sept. 1988, pp. 233–260.

[34] Martins, J. R. R. A., Alonso, J. J., and Reuther, J. J., "High-Fidelity Aerostructural Design Optimization of a Supersonic Business Jet," *Journal of Aircraft*, Vol. 41, No. 3, May 2004, pp. 523–530. doi:10.2514/1.11478.

[35] Lyu, Z., Kenway, G. K., and Martins, J. R. R. A., "Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark," *AIAA Journal*, Vol. 53, No. 4, April 2015, pp. 968–985. doi:10.2514/1.J053318.

[36] Le Pape, A. and Beaumier, P., "Numerical optimization of helicopter rotor aerodynamic performance in hover," *Aerospace science and technology*, Vol. 9, No. 3, 2005, pp. 191–201.

[37] Nielsen, E. J., Lee-Rausch, E., and Jones, W. T., "Adjoint-based design of rotors in a noninertial reference frame," *Journal of Aircraft*, Vol. 47, No. 2, 2010, pp. 638–646.

[38] Dumont, A., Le Pape, A., Peter, J., and Huberson, S., "Aerodynamic shape optimization of hovering rotors using a discrete adjoint of the Reynolds-Averaged Navier–Stokes Equations," *Journal of the American Helicopter Society*, Vol. 56, No. 3, 2011, pp. 1–11.

[39] Choi, S., Lee, K., Potsdam, M. M., and Alonso, J. J., "Helicopter rotor design using a time-spectral and adjoint-based method," *Journal of Aircraft*, Vol. 51, No. 2, 2014, pp. 412–423.

[40] Allen, C. B. and Rendall, T. C. S., "CFD-based optimization of hovering rotors using radial basis functions for shape parameterization and mesh deformation," *Optimization and Engineering*, Vol. 14, No. 1, 2013, pp. 97–118.

[41] Fingersh, L. J., Simms, D., Hand, M., Jager, D., Cotrell, J., Robinson, M., Schreck, S., and Larwood, S., "Wind tunnel testing of NREL's unsteady aerodynamics experiment," *AIAA paper*, Vol. 35, 2001, pp. 1–7.

[42] Sederberg, T. W. and Parry, S. R., "Free-form Deformation of Solid Geometric Models," *SIGGRAPH Comput. Graph.*, Vol. 20, No. 4, Aug. 1986, pp. 151–160. doi:10.1145/15886.15903.

[43] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., "A CAD-Free Approach to High-Fidelity Aerostructural Optimization," *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, Fort Worth, TX, Sept. 2010. doi:10.2514/6.2010-9231, AIAA 2010-9231.

[44] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, January 2014, pp. 144–160. doi:10.2514/1.C032150.

[45] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J. J., "Unsteady Turbomachinery Computations Using Massively Parallel Platforms," *Proceedings of the 44th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2006, AIAA 2006-0421.

[46] Lyu, Z., Kenway, G. K., Paige, C., and Martins, J. R. R. A., "Automatic Differentiation Adjoint of the Reynolds-Averaged Navier–Stokes Equations with a Turbulence Model," *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, Jul. 2013. doi:10.2514/6.2013-2581.

[47] Mader, C. A. and Martins, J. R. R. A., "Derivatives for Time-Spectral Computational Fluid Dynamics Using an Automatic Differentiation Adjoint," *AIAA Journal*, Vol. 50, No. 12, December 2012, pp. 2809–2819. doi:10.2514/1.J051658.

[48] Roe, P. L., "Approximate Riemann solvers, parameter vectors, and difference schemes," *Journal of computational physics*, Vol. 43, No. 2, 1981, pp. 357–372.

[49] Jameson, A., Schmidt, W., and Turkel, E., "Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes," *AIAA Paper*, Vol. 1259, 1981, pp. 1–19.

[50] Spalart, P. R. and Allmaras, S. R., "A one-equation turbulence model for aerodynamic flows," *AIAA Journal*, Vol. 92, No. 0439, 1992, pp. 1–22.

[51] Martins, J. R. R. A., Sturdza, P., and Alonso, J. J., "The Complex-Step Derivative Approximation," *ACM Transactions on Mathematical Software*, Vol. 29, No. 3, September 2003, pp. 245–262. doi:10.1145/838250.838251.

[52] Martins, J. R. R. A. and Hwang, J. T., "Review and Unification of Methods for Computing Derivatives of Multidisciplinary Computational Models," *AIAA Journal*, Vol. 51, No. 11, November 2013, pp. 2582–2599. doi:10.2514/1.J052184.

[53] Peter, J. E. V. and Dwight, R. P., "Numerical Sensitivity Analysis for Aerodynamic Optimization: A Survey of Approaches," *Computers and Fluids*, Vol. 39, No. 3, March 2010, pp. 373–391. doi:10.1016/j.compfluid.2009.09.013.

[54] Nadarajah, S. and Jameson, A., "A Comparison of the Continuous and Discrete Adjoint Approach to Automatic Aerodynamic Optimization," *Proceedings of the 38th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2000, AIAA 2000-0667.

[55] Hascoet, L. and Pascual, V., "The Tapenade automatic differentiation tool: Principles, model, and specification," *ACM Trans. Math. Softw.*, Vol. 39, No. 3, May 2013, pp. 20:1–20:43. doi:http://dx.doi.org/10.1145/2450153.2450158.

[56] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM journal on optimization*, Vol. 12, No. 4, 2002, pp. 979–1006.

[57] Gill, P. E., Murray, W., and Saunders, M. A., *User's Guide for SNOPT Version 7, A Fortran Package for Large-Scale Nonlinear Programming*, Systems Optimization Laboratory, Stanford University, Stanford, CA 94305, Sept. 2005.

[58] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., "pyOpt: A Python-Based Object-Oriented Framework for Nonlinear Constrained Optimization," *Structural and Multidisciplinary Optimization*, Vol. 45, No. 1, January 2012, pp. 101–118. doi:10.1007/s00158-011-0666-3.

[59] Hand, M. M., Simms, D. A., Fingersh, L., Jager, D. W., Cotrell, J. R., Schreck, S., and Larwood, S. M., *Unsteady aerodynamics experiment phase VI: wind tunnel test configurations and available data campaigns*, National Renewable Energy Laboratory, Golden, CO, 2001.

[60] Simms, D. A., Schreck, S., Hand, M., and Fingersh, L. J., *NREL unsteady aerodynamics experiment in the NASA-Ames wind tunnel: a comparison of predictions to measurements*, National Renewable Energy Laboratory, 1617 Cole Boulevard Golden, Colorado USA, 2001.

[61] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., "Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Derivative Computations," *AIAA Journal*, Vol. 52, No. 5, May 2014, pp. 935–951. doi:10.2514/1.J052255.

[62] Ashuri, T., Zaaijer, M. B., Martins, J. R. R. A., van Bussel, G. J. W., and van Kuik, G. A. M., "Multidisciplinary Design Optimization of Offshore Wind Turbines for Minimum Levelized Cost of Energy," *Renewable Energy*, Vol. 68, August 2014, pp. 893–905. doi:10.1016/j.renene.2014.02.045.